A metric for sets of trajectories that is practical and mathematically consistent

José Bento Email: jose.bento@bc.edu

arXiv:1601.03094v1 [cs.CV] 12 Jan 2016

***Abstract*—Metrics on the space of sets of trajectories are important for scientists in the field of computer vision, machine learning, robotics and general artificial intelligence. Yet existing notions of closeness are either mathematically inconsistent or of limited practical use. In this paper we outline the limitations in the existing mathematically-consistent metrics, which are based on [1], and the inconsistencies in the heuristic notions of closeness used in practice, whose main ideas are common to the CLEAR MOT measures [2] widely used in computer vision. In two steps we then propose a new intuitive metric between sets of trajectories and address these problems. First we explain a natural solution that leads to a metric that is hard to compute. Then we modify this formulation to obtain a metric that is easy to compute and keeps all the good properties of the previous metric. In particular, our notion of closeness is the first that has the following three properties: it can be quickly computed, it incorporates confusion of trajectories' dentity in an optimal way and it is a metric in the mathematical sense.**

I. Introduction

Consider two sets A and B each with multiple trajectories. We denote the set of all trajectories as $\mathcal{T}$ and the set of all finite sets of trajectories as $\mathcal{S}$. For concreteness, we assume that $A = \{A_1, ..., A_k\} \in \mathcal{S}$ and $B = \{B_1, ..., B_l\} \in \mathcal{S}$ where each trajectory, say A_1, is a set of state-time pairs $\{(t_1, x_1), ..., (t_n, x_n)\} \in \mathcal{T}$ with $t_i \in \mathbb{N}$ and $x_i \in \mathbb{R}^p$, which we can also represent as $x_i = A_1(t_i)$. In this paper we propose a new distance measure $\mathcal{D}(A, B)$ for $A, B \in \mathcal{S}$.

Measures of distance between sets of trajectories are common, for example, in the field of computer vision tracking. In this case A contains the trajectories of several objects and B is the set of approximate trajectories that a tracking algorithm reproduces from video-data. A state-time pair (t, x) can be, for example, the position and velocity of an object at time t. $\mathcal{D}(A, B)$ is important to evaluate the performance of multi-object tracking algorithms and to distinguish a good tracker, $\mathcal{D}(A, B) =$ "small", from a bad, tracker $\mathcal{D}(A, B) =$ "large".

Machine learning is another example where a measure of closeness between sets of trajectories is important. Imagine that each datapoint in a data set is a set of the trajectories followed by all players in a team during a football play. Now consider the following three tasks: the unsupervised task of clustering football plays, the supervised task of building a classifier for the different plays and the information retrieval task of finding plays similar to a reference play. We can solve these tasks using algorithms that only require a measure of distance on the sets of trajectories, e.g. [3], [4], [5] respectively.

Our goal is that $\mathcal{D}(A, B)$ is both mathematically consistent and useful in practice. A common mathematical inconsistency of similarity measures used in practical applications involving trajectories is that they are not a mathematical metric. In Section III we show this is the case with the CLEAR MOT measures [2], widely used in computer vision, and hence also with many other distance measures that use heuristics similar to the CLEAR MOT. On the contrary, the measure we introduce satisfies the properties of a metric: for every $A, B, C \in \mathcal{S}$ we have (i) **(non-negativity)** $\mathcal{D}(A, B) \geq 0$, (ii) **(coincidence)** $\mathcal{D}(A, B) = 0$ iff $A = B$, (iii) **(symmetry)** $\mathcal{D}(A, B) = \mathcal{D}(B, A)$ and (iv) **(sub-additivity)** $\mathcal{D}(A, C) \leq \mathcal{D}(A, B) + \mathcal{D}(B, C)$.

In the context of tracking, it is easy to see why not dealing with a mathematical metric can lead to inconsistencies. Imagine that under a certain distance $\mathcal{D}$ both tracker1's output, O_1, and tracker2's output, O_2, are close to the ground truth, GT. In this setting, one intuitively expects that O_1 and O_2 are also close to each other. This expectation is related to the sub-additivity property $\mathcal{D}(O_1, O_2) \leq \mathcal{D}(O_1, GT) + \mathcal{D}(O_2, GT)$ and we consider its violation an inconsistency. The need for symmetry, non-negativity and the coincidence property is intuitive. Note that the machine learning algorithms mentioned above, namely [3], [4], [5], also require $\mathcal{D}$ to be a metric.

In addition to seeking mathematical consistency, our goal is also to define a useful measure. In the examples that follow we again focus on vision tracking because in this context it is clear why the existing mathematically consistent measures, which are based on [1], are not useful in practice.

We obviously want $\mathcal{D}(A, B)$ to have a computation time that scales well with the number and length of the trajectories in A and B. We show our metric is fast to compute in Section V-A. In addition, to be useful, $\mathcal{D}(A, B)$ must behave as expected in scenarios with intuitive answers as, for example, in the scenarios we now explain. These help us understand how a consistent and practical $\mathcal{D}$ should look like, help us introduce relevant previous work, some notation, and also help understanding why our task is not easy.

Scenario 1: In Fig. 1-(a) the set A has only one trajectory $\{(1, -1.0), (2, -0.6), (3, -0.2), (4, 0.2), (5, 0.6), (6, 1.0)\}$ and the set B has only one trajectory $\{(1, -0.9), (2, -0.54), (3, -0.18), (4, 0.18), (5, 0.54), (6, 0.9)\}$. One intuitive way to define $\mathcal{D}(A, B)$ in this case is to form the vectors $[-1.0, -0.6, -0.2, 0.2, 0.6, 1.0]$ and $[-0.9, -0.54, -0.18, 0.18, 0.54, 0.9]$ and compute the sum of the distances between each of their components. This gives $\mathcal{D}(A, B) = 0.36$. More generally, we can use any vector norm between these two vectors and, as long as both A and B only have one trajectory of equal length, this procedure defines a metric,

$$\mathcal{D}(A, B) = \|A - B\|,$$

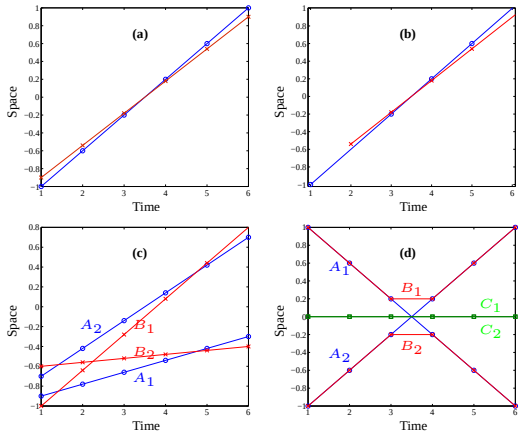

Fig. 1. Scenarios where defining $\mathcal{D}(A, B)$ is intuitive. The red cross markers $\times$ are the trajectories from B, the blue circle markers $\circ$ are the trajectories from A and the green square markers $\square$ from C.

where A and B on the right hand side are the vector representation of the trajectories A and B.

Scenario 2: In Fig. 1-(b) the sets A and B still only have one trajectory but now the trajectories have "holes". To define $\mathcal{D}$ for this scenario we can reduce it to Scenario 1 by computing the total distance between each of the components of the vectors $[-1.0, *, -0.2, 0.2, 0.6, *]$ and $[*, -0.54, -0.18, 0.18, 0.54, *]$ and using the rule that $|x - *| = M > 0$ if $x \in \mathbb{R}$ and that $|* - *| = 0$. For example, if $M = 0.1$ we get $\mathcal{D}(A, B) = 0.3$. We can interpret the symbol * as total uncertainty in a tracker A regarding the position of an object B at a particular point in time, in which case we penalize its performance by a fixed amount. If tracker A outputs * because the object B does not exist in that time instant (e.g. it is outside the tracker's range) we do not penalize its performance, hence $|* - *| = 0$.

More generally, let $A^+ = [..., A^+(t), ...]$ and $B^+ = [..., B^+(t), ...]$ be the vector representation of the single trajectories in A and B, now with each point living in $\mathbb{R}^p$, and padded with the symbol $*$ and let $d^+ : \mathbb{R}^p \cup \{*\} \times \mathbb{R}^p \cup \{*\} \mapsto \mathbb{R}$ be an extension of some point-wise metric $d : \mathbb{R}^p \times \mathbb{R}^p \mapsto \mathbb{R}$ such that for every $x, y \in \mathbb{R}$ we have $d^+(x, y) = \min\{2M, d(x, y)\}$, $d^+(x, *) = d^+(*, x) = M > 0$ and $d^+(*, *) = 0$. One can show that d^+ is a metric (c.f. Lemma 2 in Appendix B) and hence so is

$$\mathcal{D}(A, B) = \sum_t d^+(A^+(t), B^+(t)).$$

Scenario 3: For Fig. 1-(c) we cannot define $\mathcal{D}(A, B)$ as easily as in Scenario 1 and 2 because A and B have multiple trajectories. We can try the same approach and reduce $\mathcal{D}$ to a norm in a vector space by (a) writing A as the concatenation of $A_1 = [-0.90, -0.78, -0.66, -0.54, -0.42, -0.30]$ and $A_2 = [-0.7, -0.42, -0.14, 0.14, 0.42, 0.70]$ and (b) writing B as the concatenation of $B_1 = [-1.00, -0.64, -0.28, 0.08, 0.44, 0.80]$ and $B_2 = [-0.60, -0.56, -0.52, -0.48, -0.44, -0.40]$ and (c) computing the total distance between the components of the vectors A and B thus formed. However, this procedure is not well defined because if we concatenate $[A_1, A_2]$ and $[B_1, B_2]$ we get $\mathcal{D}(A, B) = 6.40$ and if we concatenate $[A_2, A_1]$ and $[B_1, B_2]$ we get $\mathcal{D}(A, B) = 1.68$. Effectively we have defined $\mathcal{D}(A, B)$ by comparing each trajectory in A against one trajectory in B but since A and B are unordered sets we do not know if we should compare A_1 with B_1 or with B_2.

One way to resolve this ambiguity is to choose the smallest of the values computed, 1.68 in our example. That is, we choose the most favorable comparison between trajectories. The authors in [6] show that this procedure results in a metric. In this paper we call this metric OSPA because it is basically an extension of the OSPA metric of [1] from A and B being sets of points to A and B being sets of trajectories.

To write a general expression for OSPA we first extend A and B such that each have m trajectories of length T. First we make all trajectories have the same length T by padding them with state-time pairs $(*, t)$. Padded trajectories are represented by A_i^+ and B_i^+. Afterwards, if A has k trajectories and B has l trajectories we extend A to $A^+ = \{A_1^+, ..., A_k^+, A_{k+1}^+, ..., A_m^+\}$ and we extend B to $B^+ = \{B_1^+, ..., B_l^+, B_{l+1}^+, ..., B_m^+\}$ where $m = k + l$ and where the extra trajectories A_i^+, $i > k$, and B_i^+, $i > l$, are all of the kind $\{(1, *), ..., (T, *)\}$. We write the set of all extended trajectories as $\mathcal{T}^+$. Let Π be the set of all permutations of $[m] \equiv \{1, ..., m\}$. We represent any $\sigma \in \Pi$ as an ordered set $(\sigma_1, ..., \sigma_m)$ where $\sigma_i \in [m]$.

Definition 1. *The OSPA metric is defined as*

$$\mathcal{D}(A, B) = \min_{\sigma \in \Pi} \sum_{i,t} d^+(A_i^+(t), B_{\sigma_i}^+(t)).$$

This metric requires a procedure to match trajectories in A with trajectories in B for comparison and matching (parts of) trajectories is also a central theme in our paper. Indeed, early on researchers identified finding such a match as a central task in defining a metric for sets of trajectories [7], [8].

In the context of computer vision tracking it is important that A and B are extended to have have $m = l + k$ trajectories each. Let A be the ground-truth trajectories and B be the tracker's reconstructed trajectories. We can match any of the k original trajectories in A to a trajectory $B_i^+ = \{(1, *), ..., (T, *)\}$, $i > l$, to represent that the tracker missed a trajectory completely and we can match any of the l original trajectories in B to a trajectory $A_i^+ = \{(1, *), ..., (T, *)\}$, $i > k$, to represent the tracker producing a spurious trajectory.

The next scenario shows why OSPA-based metrics can be uninformative in practice.

Scenario 4: For the scenario in Fig. 1-(d), we cannot use an OSPA-based metric to obtain a $\mathcal{D}(A, B)$ that is practical in applications such as performance evaluation in computer vision tracking. Indeed, if we use a procedure that matches full-trajectories to full-trajectories as described in Scenario 3, and after we compute the sum of the absolute values of the difference between components of the matched trajectories, we obtain a value of $\mathcal{D}(A, B) = 7.2$. We obtain this value regardless of whether we compare $A_1 = \{(1, 1.00), (2, 0.60), (3, 0.20), (4, -0.20), (5, -0.60), (6, -1.00)\}$ with $B_1 = \{(1, 1.00), (2, 0.60), (3, 0.20), (4, 0.20), (5, 0.60)$ and $A_2 = \{(1, -1.00), (2, -0.60), (3, -0.20), (4, 0.20), (5, 0.60), (6, 1.00)\}$ with $B_2 = \{(1, -1.00), (2, -0.60), (3, -0.20), (4, -0.20), (5, -0.60), (6, -1.00)\}$, or we compare A_1 with

B_2 and A_2 with B_1. Note however that if we compute $\mathcal{D}(A,C)$ or $\mathcal{D}(B,C)$, where C is an object that is stopped at position 0, we also get 7.2. If we interpret A as the trajectories of two objects and B and C as the output of two different trackers, the OSPA-based metric is saying that the trackers are equally good, yet, most people would probably resist equating both because the tracker that outputs B is actually doing a good job at estimating positions but simply makes one confusion of identity of the objects A_1 and A_2 as they pass by each while the other tracker miss estimates the objects' positions all the time.

OSPA-based metrics tend to over penalize the quality of a tracker in all situations like in Fig. 1-(d) where one trajectory in A follows closely one trajectory in B but later follows another trajectory in$B$. When this happens we say we have an **identity switch** and the fact that OSPA-based metrics cannot deal with identity switches makes then useless in many important applications. For example, imagine that we are tracking hundreds of bats, just like the authors [9], [10], [11] did, and we use a tracker that has the impressive ability to simultaneously track all the bats flying around a cave very accurately for many minutes. Only once in a while the tracker confuses the identity of two bats that pass by very close to each other. According to OSPA, the performance of this tracker is very poor. However, probably most people agree that its performance is above that of most existing trackers.

It is because of examples as the above that the computer vision tracking community mostly uses measures of closeness like the CLEAR MOT to assess tracking performance. The CLEAR MOT are in fact two measure of closeness, MOTP and MOTA. Although they are not a metric, they do allow to distinguish good trackers from the bad trackers in realistic scenarios. In particular, they allow the user to control how much identity switches should penalize the value of $\mathcal{D}(A,B)$.

The main idea behind the CLEAR MOT is to use a simple heuristic to match *parts* of trajectories in A *to parts* of trajectories in B instead of matching full trajectories in A to full trajectories in B like in the OSPA-base metrics. Let Π^T be the set of all sequences of length T of permutations of $[m]$ objects. In other words, if $\Sigma \in \Pi^T$ then $\Sigma = (\sigma(1), \sigma(2), ..., \sigma(T))$ where $\sigma(t) \in \Pi, \forall t$. The CLEAR MOT heuristic sequentially builds a sequence $\Sigma_{\text{MOT}} = (\sigma_{\text{MOT}}(1), \sigma_{\text{MOT}}(2), ..., \sigma_{\text{MOT}}(T)) \in \Pi^T$ in such a way that $\mathcal{D}(A,B)$, as defined in Definition 2, is small and at the same time the number of times $\sigma(t)$ is different from $\sigma(t+1)$ is also small.

Definition 2. *The CLEAR MOT's distance measure for evaluating tracking precision is called MOTP and is defined as*

$$\mathcal{D}(A,B) = \sum_{i,t} d^+(A_i^+(t), B^+_{\sigma_{MOT_i}(t)}(t)).$$

The heuristic focuses on simplicity to guarantee that we can compute $\mathcal{D}(A,B)$ fast. Unfortunately, at the expense of being computable, it inherits serious limitations as we explain in Section III. It also depends on a user-defined parameter thr that controls how easily $\sigma_{\text{MOT}}(t)$ changes from one time instant to the other. Intuitively, thr allows the user to say how close two objects must be for MOTP to have confidence that they should be compared.

Definition 3. *The CLEAR MOT matching heuristic defines Σ_{MOT} sequentially as follows.*

1) *Initialize $\sigma_{MOT}(1)$ such that $\sum_i d^+(A_i^+(1), B^+_{\sigma_{MOT_i}(1)}(1))$ is minimal;*
2) *For each $t > 1$ do: for all $i, j \in [m]$ such that $\sigma_{MOT_i}(t-1) = j$ and $d^+(A_i^+(t), B_j^+(t)) < thr$ fix $\sigma_{MOT_i}(t) = j$. We call such matches as* **anchored**. *Define the non-fixed components of $\sigma_{MOT}(t)$ such that $\sum_i d^+(A_i^+(t), B^+_{\sigma_{MOT_i}(t)}(t))$ is minimal.*

What is the value of MOTP for an example like in Fig. 1-(d) if $thr = 0.19$? First we need to compute Σ_{MOT}. Let us focus on the pair A and B. For $t = 1$ we minimize distance by matching A_1 to B_1 and A_2 to B_2. This is mathematically equivalent to setting $\sigma_{\text{MOT}}(1) = (1, 2)$. For $t = 2$ the previous association is anchored because $d^+(A_1^+(1), B^+_{\sigma_{\text{MOT}_1}(1)}(1)) = d^+(A_2^+(1), B^+_{\sigma_{\text{MOT}_2}(1)}(1)) = 0 < thr$ and so $\sigma_{\text{MOT}}(2) = (1, 2)$. The same happens for $t = 3$. For $t = 4$ the previous association cannot be anchored because that would lead to distances larger than thr so we need to re-minimize the distance. This leads to $\sigma_{\text{MOT}}(4) = (2, 1)$. For $t = 5$ and $t = 6$ this association is anchored. In short, we get $\Sigma_{\text{MOT}} = ((1,2),(1,2),(1,2),(2,1),(2,1),(2,1))$. Using this value we compute $\mathcal{D}(A,B) = 0$. We also get $\mathcal{D}(A,C) = \mathcal{D}(B,C) = 7.2$. In other words, MOTP says that, apart from an identity switch, tracker B is better than tracker C, which is the intuitive answer we expect.

Recall that, as we said before, although MOTP defines a measure of closeness between sets of trajectories that is useful, it is mathematically inconsistent. In other words, there is still no measure that is both mathematically consistent and useful in measuring the distance between sets of multiple trajectories. Our main contribution is to show how to capture all the good properties of all the definitions above, and none of their limitations, by using a definition of the form

$$\mathcal{D}(A,B) = \min_{\Sigma \in \Pi^T} \left\{ \mathcal{K}(\Sigma) + \sum_{t=1}^{T} \sum_{i=1}^{m} d^+(A_i^+(t), B^+_{\sigma_i(t)}(t)) \right\} \quad (1)$$

where $\mathcal{K} : \Pi^T \mapsto \mathbb{R}$ is a function that penalizes $\sigma(t)$ changing with time. With this we fill an important gap in the literature,

More comprehensively, our contribution is the following:

1) We show that the CLEAR MOT association heuristic can be unintuitive;
2) We show that MOTP does not define a mathematical metric;
3) We use (1) to define a similarity measure that we prove is a metric and that deals with identity switches optimally;
4) We modify (1) to define the first similarity measure that is fast to compute, is a metric, and is useful in practice;

We now present each of our contributions in order. We start however with an extensive literature review to clearly delineate the gap we will fill.

Note that in this paper we repeatedly make reference to the CLEAR MOT measures and OSPA-based metrics. Not because there are no other metrics beside these two but because they are excellent representatives of the two categories in which most the work done in this area falls: similarity measures that

are practical (although even these have problems) but are not a metric and similarity measures that are a metric but are not practical.

II. Related work

In this section we review existing work on metrics for sets of trajectories. Among all the work we review we find no similarity measure between two sets of trajectories A and B that is mathematically consistent (a metric) and, at the same time, is useful and can deal with identity switches (it looks for similarities between parts of trajectories of A and parts of trajectories of B). As we argue in the introduction, these are two desired characteristic that we are the first to incorporate and we focus our discussion around them. Several other ideas in the work we review can be easily incorporated into our work and define new metrics that compete with past work beyond these two characteristics. The discussion of these variant metrics is beyond the scope of this paper.

Our work is related to the problem of defining a distance between two sets A and B, a topic too vast to review in this paper. In [12] the reader can find many of these distances. In this paper, the sets A and B are sets of trajectories, and this limits the scope of our discussion. In the simplest case, when trajectories have only one point/vector, typical definitions involve computing an average of the distance between all pairs of elements from A and B, e.g. [13], or the sum of distance between a few pairs of elements from A and B obtained by some procedure that matches elements of A with elements of B, e.g. [1], [14]. In general however, each trajectory is composed by a set of vectors indexed by time, which limits our discussion even more. To the best of our knowledge the most rigorous works on metrics for sets of trajectories are based on the ideas proposed in [1]. In [1] the authors propose a distance between sets of vectors, the optimal sub-pattern assignment metric (OSPA), and explain its advantages for evaluating the performance of multi-object filters compared to other distances between sets. In particular, the OSPA metric has better sensitivity to cardinality differences between sets than the Hausdorff metric and does not lead to complicated interpretations as the optimal mass transfer metric (OMAT) that [15] propose to address the limitations of the Hausdorff metric. The OSPA metric between two sets of vectors is well defined for any metric between vectors.

All the spin-offs of [1] focus on defining a metric between two sets of trajectories A and B for the purpose of evaluating the performance of tracking algorithms. We recall however that many applications in machine learning and AI apart from tracking benefit if we work with a metric, for example like the ones we define in this paper, rather than a similarity measure that is not a metric. It is crucial to note that all the spin-offs of [1] only compare full trajectories in A to full trajectories in B and hence suffer from the same limitations that we describe in Example 4 in the Introduction.

The authors in [6] define the OSPA-T metric in two steps. In the first step they solve an optimization problem that optimally matches full tracks in A to full tracks in B while taking into account that tracks have different lengths and might be incomplete. In the second step, they assign labels to each track based on this match, they compute the OSPA metric for each time instant using a new metric between pair of vectors that considers both the vectors' components as well as their labels and they sum all the OSPA values across all time instants. Although the optimization problem of the first step defines a metric, [16] point out that the full two-step procedure that defines OSPA-T can violate the triangle inequality.

The authors in [16] define the OSPAMT to be a metric and to be more reliable than OSPA-T when we evaluate the performance of multi-target tracking algorithms. The OSPAMT metric also computes an optimal match between full trajectories in A and full trajectories in B but unlike OSPA-T allows to match one full trajectory in A to multiple full trajectories in B (and vice-versa). The authors make this design choice not to penalize a tracker when it outputs only one track for two objects that move closely together.

Some extensions to OSPA incorporate the uncertainty in the measurements. The Q-OSPA metric defined in [17] incorporates uncertainty by weighting the distance between pairs of points by the product of their certainty and by adding a new term that is proportional to the product of the uncertainties. The H-OSPA metric defined in [18] incorporates uncertainty by using OSAP with distributions as elements of A and B instead of vectors and using the Hellinger distance between distributions instead of the Euclidean distance between vectors. The authors of both works focus only on the simpler case where the sets A and B contain vectors/points and not trajectories. However, combining their work with that of [6] or [16] to obtain a metric between sets of trajectories is immediate.

The papers above are fairly recent and the search for similarity measures between sets of trajectories that are a metric is not older. However, researchers in the field of computer vision have been interested in defining similarly measures for sets of trajectories to evaluate the performance of tracking algorithms much prior to these works. It is impossible to review all work done in this area. Specially because the evaluation of the performance of trackers has many challenges other than the problem of defining a similarity measure. See [19], [20] for some examples of these other challenges. Nonetheless, we mention a few works and point out ideas in them that relate to our problem. We emphasize that none of the following works defines a metric mathematically.

Most of our paper is about solving the problem of forming a good match between elements in A and B. In [7], which is expanded in a subsequent paper [8], the authors are one of the first to identify this as the central problem in defining a similarity measure, although some of their ideas draw from the much earlier Ph.D. thesis of one of the authors [21]. They propose a one-to-one association between the different points of A and B but this association is optimally computed independently every instant and there is no discussion about the number of changes in matching that this might create.

One of the reasons why the CLEAR MOT metrics are widely used is because they create a simple association between A and B, i.e., the association between A and B does not change often in time. It appears that [22] is one of the first works that describes how to control the number of association

changes when computing a similarity measure. The authors do not associate A and B independently every time instant but rather use a sequential matching procedure that tries to keep the association from the previous time instant if possible. This is similar to the procedure used in the much more recent CLEAR MOT measures that we discuss in the Introduction.

The association that [22] use at every point in time is not one-to-one optimal like in [7] or in the CLEAR MOT, rather the authors use a simple thresholding rule to associate neighboring elements of A and B. The idea of using a simple threshold rule to compare A and B seems to have survived until relatively recent. For example, in [23] the authors match a full trajectory in A to a full trajectory in B if they are close in space for a sufficiently long time interval. The authors in [24] use a similar thresholding method to match A and B.

Shortly after the papers [7], [8] some of the same authors discuss again the problem of associating A and B. In [25] they propose four different methods to solve this problem. In the first method, A and B are optimally matched independently at each time instant, possibly generating different associations every time instant. In the second method, the association between A and B cannot change with time. This is similar to what happens in the OSPA-based metrics. It creates simple associations but can also lead to problems as we discuss in the Introduction. In the third method, the authors allow the association between A and B to change in time but only in special circumstances. Unfortunately, as they point out, this leads to a NP-hard problem. Finally, in the fourth method, the authors discuss some ad-hoc ideas to try to minimize the number of changes in association. Some of these ideas resemble the idea used in the CLEAR MOT of keeping the association from the previous time period if possible.

Not all challenges in defining a measure of performance for trackers are related to defining a similarity measure between trajectories. In addition, not all challenges in defining a similarity measure between trajectories are related to defining a way to compare and associate the trajectories in A with the trajectories in B. Even if we assume A and B are already matched, there is the question of computing some quantity from this match. A typical quantity researchers compute is the sum of the distances between the different points of matched trajectories, averaged over all pairs of matched trajectories. As explained later, our metric is a combination of this quantity with the number of confusion of identities in the match between A and B.

Apart from other distance-related quantities, researchers also compute many quantities that evaluate the quality of the match itself. For example, the number of trajectories in A that do not match to any trajectory in B or the number of times in which the match between A and B changes. [26], [27] list many quantities that can be computed to measure the similarity between A and B and some of these quantities might be useful to define similarity measures for applications outside tracking.

It is worth mentioning a few works that differ from the main stream in this aspect. One of them is [28], where the authors define a similarity measure based on comparing the occurrence of special discrete events in A and B, and another is [29], where the authors propose an information theoretic measure of similarity between sets of trajectories. Finally, in [30] the authors propose a similarity measured based on hidden Markov models that does not assume that the temporal sampling rates of the trajectories are equal.

III. The limitations of the CLEAR MOT

The CLEAR MOT similarity measures are one of the most commonly used measures that are able to associate parts of trajectories of A with parts of trajectories in B such that the sum of the distance between associated points is small and at the same time the association is simple, i.e. does not change often in time. The CLEAR MOT achieve this using the heuristic procedure in Definition 3. In Theorem 1 we show that this heuristic can be unintuitive and in Theorem 2 we show that MOTP is mathematically inconsistent.

Theorem 1. *For any $\Sigma \in \Pi^T$ and any $A, B \in \mathcal{S}$ without holes and with m trajectories each let $swi(\Sigma) = \frac{1}{T-1}\sum_{t=1}^{T-1} \mathbb{I}(\sigma(t) \neq \sigma(t+1))$ and $dist(\Sigma, A, B) = \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{m} \|A_i(t) - B_{\sigma_i(t)}(t)\|$.*

For any CLEAR MOT threshold $thr > 0$, there exists A, B and Σ such that $dist(\Sigma, A, B) < \mathcal{O}(1/T)$, $swi(\Sigma) = 0$ and $dist(\Sigma_{MOTP}, A, B) > m\,thr - \mathcal{O}(1/T)$).

In the context of computer vision tracking the theorem above shows that there are situations in which the average tracking performance gets arbitrarily close to optimal with time while the CLEAR MOT association Σ_{MOT} says that it does not. In other words, the CLEAR MOT can be unintuitive.

The following theorem show that, for any $thr > 0$, the MOTP is mathematically inconsistent.

Theorem 2. *MOTP is not a metric.*

The proofs of these two theorems are in Appendix A.

IV. A natural metric for sets of trajectories

Despite the large volumes of research about similarity measures for sets of trajectories, largely motivated, at least initially, by the need to assess the performance of tracking systems, our work is the first to propose a measure that is both mathematically consistent and practical.

To build towards our final metric, in this section we define a metric that *captures the best match between the parts of trajectories in A and parts of trajectories in B while tacking into account that simple matches are preferred to complex ones.* As the four scenarios in the introduction suggest, this is a desirable characteristic for a measure of distance and is only partially present in the OSPA and the CLEAR MOT which are representatives of the two currently segregated camps of (a) the consistent metrics and of (b) the practical metrics. Unfortunately the metric we now define is hard to compute and we need an extra section to modify it and reduce computational complexity.

Before we introduce our metric we introduce some new notation. We denote the identity permutation by $I = (1, 2, ..., m)$. We also define Π^T as all sequence of permutations of length T. In particular if $\Sigma \in \Pi^T$ then $\Sigma = (\sigma(1), \sigma(2), ..., \sigma(T))$ where $\sigma(t) \in \Pi, \forall t$. Given $\Sigma \in \Pi^T$, we define $\Sigma^{-1} \in \Pi^T$ as $\Sigma^{-1} \equiv (\sigma(1)^{-1}, \sigma(2)^{-1}, ..., \sigma(T)^{-1})$ where $\sigma(t)^{-1}$

is the inverse permutation of $\sigma(t)$. In addition, if $\Sigma' = (\sigma'(1), \sigma'(2), ..., \sigma'(T)) \in \Pi^T$ is another sequence of permutations we define $\Sigma \circ \Sigma' \in \Pi^T$ as $\Sigma \circ \Sigma' \equiv (\sigma(1) \circ \sigma'(1), ..., \sigma(T) \circ \sigma'(T))$ where the symbol $\circ$ denotes composition of permutations.

We now introduce a another new definition.

Definition 4. *$\mathcal{K} : \Pi^T \mapsto \mathbb{R}_0^+$ is a map that satisfies the following three properties (i) $\mathcal{K}(\Sigma) = 0$ if and only if Σ is constant $\Sigma = (\sigma, \sigma, ..., \sigma)$ for some $\sigma \in \Pi$, (ii) $\mathcal{K}(\Sigma^{-1}) = \mathcal{K}(\Sigma)$ and (iii) $\mathcal{K}(\Sigma \circ \Sigma') \leq \mathcal{K}(\Sigma) + \mathcal{K}(\Sigma')$.*

We now define a new distance measure.

Definition 5. *The* **natural distance** *between two sets of trajectories is a map $\mathcal{D}_{nat} : \mathcal{S} \times \mathcal{S} \mapsto \mathbb{R}_0^+$ such that for any $A, B \in \mathcal{S}$*

$$\mathcal{D}_{nat}(A, B) = \min_{\Sigma \in \Pi^T} \left\{ \mathcal{K}(\Sigma) + \sum_{t=1}^{T} \sum_{i=1}^{m} d^+(A_i^+(t), B_{\sigma_i(t)}^+(t)) \right\}. \quad (2)$$

Remark 1. *This definition is compatible with scaling the term $\mathcal{K}(\Sigma)$ by some constant $\alpha > 0$. We are not ignoring this possibility but simply put any scaling factor 'inside' $\mathcal{K}$. Indeed, if $\mathcal{K}(\Sigma)$ satisfies properties (i), (ii) and (iii) then so does $\alpha\mathcal{K}(\Sigma)$. It is the user's decision to determine, for the application she has in mind, the relative importance of the terms $\mathcal{K}$ and d^+ in* (2)*.*

This definition is related to the CLEAR MOT measure in the sense that it builds an association between parts of trajectories in A and parts of trajectories in B such that the distance between associated trajectories is small and the association does not change often in time. It is the term $\mathcal{K}$ in the definition of $\mathcal{D}_{nat}$ that allows us to control the complexity of the match between the trajectories in A and in B. This guarantees that $\mathcal{D}_{nat}$ is as useful in practice as the CLEAR MOT.

This definition is related to the OSPA metric regarding mathematical consistency but is much more general. Indeed, we recover the OSPA metric if we confine ourselves to $\mathcal{K} = \mathcal{K}_{OSPA}$ defined as $\mathcal{K}_{OSPA}(\Sigma) = 0$ if $\Sigma = (\sigma, \sigma, ..., \sigma)$ and $\mathcal{K}_{OSPA}(\Sigma) = \infty$ otherwise. Using $\mathcal{K} = \mathcal{K}_{OSPA}$ forces us to compare full trajectories in A to full trajectories in B, something that as we explained in the Introduction, leads to unintuitive results.

Strictly speaking our definition does not define a single measure but a family of measures,one for each choice of $d(.)$ and $\mathcal{K}(.)$. The following theorem shows that all the measures in this family are a metric. Its proof is in Appendix B.

Theorem 3. *The map $\mathcal{D}_{nat}$ is a metric on $\mathcal{S}$.*

A natural choice for d is the Euclidian metric, that is, $d(x, x') = \|x - x'\|_2$. We now show that several desirable definitions for $\mathcal{K}$ satisfy the properties in Definition 4. In the context of computer vision tracking, all these $\mathcal{K}$ allow to penalize trackers that generate many identity switches.

Perhaps the most obvious definition is for $\mathcal{K}$ to count the number of times that the association between A and B changes.

Theorem 4. *Let $\mathcal{K}_{count}(\Sigma) = \alpha \sum_{t=1}^{T-1} \mathbb{I}(\sigma(t+1) \circ \sigma(t)^{-1} \neq I)$, $\alpha > 0$. $\mathcal{K}_{count}$ satisfies Definition 4.*

Another two desirable choices for $\mathcal{K}$ are (a) the function that sums the minimum number of transpositions to go from one permutation to the next and (b) the function that sums the number of adjacent transpositions to go from one permutation to the next. In what follows $k_{Cayley}(\sigma)$ gives the minimum number of transpositions to obtain the identity permutation from $\sigma \in \Pi$ and $k_{Kendall}(\sigma)$ gives the number of adjacent transposition that the bubble-sort algorithm performs when sorting σ to obtain the identity permutation. The Cayley distance can be traced back to [31] and the Kendall distance to [32].

Theorem 5. *Let $\mathcal{K}_{trans}(\Sigma) = \alpha \sum_{t=1}^{T-1} k_{Cayley}(\sigma(t+1) \circ \sigma(t)^{-1})$, $\alpha > 0$. $\mathcal{K}_{trans}$ satisfies Definition 4.*

Theorem 6. *Let $\mathcal{K}_{adjtrans}(\Sigma) = \alpha \sum_{t=1}^{T-1} k_{Kendall}(\sigma(t+1) \circ \sigma(t)^{-1})$, $\alpha > 0$. $\mathcal{K}_{adjtrans}$ satisfies Definition 4.*

The proof of Theorems 4, 5 and 6 is in Appendix B.

Although many intuitive choices for $\mathcal{K}$ satisfy properties (i), (ii) and (iii) some natural ones to not. For example, given a $\beta \geq 1$ we might want to define $\mathcal{K}_{maxcount}$ as $\mathcal{K}_{maxcount}(\Sigma) = \mathcal{K}_{count}(\Sigma)$ if $\mathcal{K}_{count}(\Sigma) \leq \beta$ and $\mathcal{K}_{maxcount}(\Sigma) = \infty$ if $\mathcal{K}_{count}(\Sigma) > \beta$ (we can replace ∞ by some very large number if we want to be technical about the range of $\mathcal{K}$ being $\mathbb{R}_0^+$). This $\mathcal{K}$ basically forces us not to create an association between A and B more complex than a certain amount β, something natural to desire. The following is proved in Appendix B.

Theorem 7. *$\mathcal{K}_{maxcount}$ does not satisfy (iii).*

Even if a function $\mathcal{K}$ violates some of the properties it might still be the case that using that function in (2) defines a metric. However, we often find that from a set of associations Σ, Σ' that violate some of the three properties in Definition 4 we can build three set of trajectories A, B and C that violate some of the properties required of a metric. This is the case, for example, for $\mathcal{K}_{maxcount}$ and d the Euclidean metric.

Theorem 8. *If $\mathcal{K} = \mathcal{K}_{maxcount}$ and d is the Euclidean distance then equation* (2) *does not define a metric.*

The proof of this theorem is in Appendix B. In this sense, the conditions in Definition 4 are both necessary and sufficient for equation (2) to be a metric.

Some choices for d and $\mathcal{K}$ lead to a metric $\mathcal{D}_{nat}$ that is easy to compute $\mathcal{D}_{nat}$. For example, if we choose $\mathcal{K} = \mathcal{K}_{OSPA}$ and d to be the Euclidean distance then $\mathcal{D}_{nat}$ reduces to OSPA metric which can be solved in polynomial time using, for example, the Hungarian algorithm [33] (cf. [1]). Unfortunately, many desirable choices for $\mathcal{K}$ and d lead to a hard problem. For example, if d is the Euclidean distance and $\mathcal{K} = \mathcal{K}_{count}$ or $\mathcal{K} = \mathcal{K}_{Cayley}$ or $\mathcal{K} = \mathcal{K}_{Kendall}$ then computing $\mathcal{D}_{nat}$ is related to solving a multi-dimensional assignment problem that is HP-hard [34].

We solve this issue in the next section where we modify $\mathcal{D}_{nat}$ to obtain a new metric that is as easy to compute as solving a linear program (recall that all LPs can be solved in polynomial time [35]).

V. A NATURAL AND COMPUTABLE METRIC

First we introduce some notation. Given $A, B \in \mathcal{S}$ we define A^+ and B^+ just like in Section IV. We define T and m accordingly. Given d we also define d^+ like in Section IV. Let $\mathcal{P}$ be the set of all doubly stochastic matrices, that is, $\mathcal{P} = \{W \in \mathbb{R}^{m\times m} : W^\dagger \mathbf{1} = \mathbf{1}, W\mathbf{1} = \mathbf{1}, 0 \leq W \leq 1\}$. Let $\mathcal{P}^T$ be the set of all sequences of length T of doubly stochastic matrices. Given $A, B \in \mathcal{S}$ and $t \leq T$ we define the matrix $D^{AB}(t) \in \mathbb{R}^{m\times m}$ as $(D^{AB}(t))_{ij} = d^+(A_i^+(t), B_j^+(t))$. Let $|||$ be a matrix norm that satisfies the following property for any four matrices $X_1, X_2, Y_1, Y_2 \in \mathcal{P}$

$$\|Y_2X_2 - X_2X_1\| \leq \|Y_2 - Y_1\| + \|X_2 - X_1\|. \quad (3)$$

Note that this is analougous to property (iii) that we use for $\mathcal{K}$ in Section IV.

Definition 6. *The* **natural computable distance** *between two sets of trajectories is a map $\mathcal{D}_{comp} : \mathcal{S} \times \mathcal{S} \mapsto \mathbb{R}_0^+$ such that for any $A, B \in \mathcal{S}$*

$$\mathcal{D}_{comp}(A, B) = \min_{\{W(t)\}\in\mathcal{P}^T} \Big\{ \sum_{t=1}^{T-1} \|W(t+1) - W(t)\| + \sum_{t=1}^{T} \mathbf{tr}(W(t)^\dagger D^{AB}(t)) \Big\}. \quad (4)$$

Remark 2. *This definition is compatible with scaling the first term by some constant $\alpha > 0$. We are not ignoring this possibility but simply put any scaling factor "inside" the matrix norm. In Section VI we solve* (4) *for different values of α to obtain a trade-off plot between its two terms.*

The motivation for this definition is to obtain a measure that is as similar as possible to $\mathcal{D}_{nat}$ but that we can compute fast. Just like MOTP and $\mathcal{D}_{nat}$, the measure $\mathcal{D}_{comp}$ also builds an association between parts of trajectories in A and parts of trajectories in B. This association is represented by the matrices $W(t)$. However, unlike Σ in $\mathcal{D}_{nat}$, the W matrices also give us a weight/strength for the different matches between A and B. Just like $\mathcal{D}_{nat}$, and as the next theorem shows, $\mathcal{D}_{comp}$ is also mathematically consistent.

Theorem 9. *The map $\mathcal{D}_{comp}$ is a metric on $\mathcal{S}$.*

The proof of this theorem is in Appendix C. Most importantly, unlike $\mathcal{D}_{nat}$, computing $\mathcal{D}_{comp}$ amounts to solving a convex optimization problem. In short, $\mathcal{D}_{comp}$ combines all the advantages of OSPA, MOTP and $\mathcal{D}_{comp}$.

The following Lemma gives sufficient conditions for a metric $\|.\|$ to satisfy property (3). See Appendix C for a proof.

Lema 1. *If $\|.\|$ is a sub-multiplicative norm and $\|W\| \leq 1$ for all $W \in \mathcal{P}$ then $\|.\|$ satisfies property* (3)*.*

This lemma implies, for example, that the 1-norm and ∞-norm spectral norm for matrices are valid choices for $\|.\|$. In particular, the use of $\|.\|_1$ in $\mathcal{D}_{comp}$ is extremely useful because it induces the changes of association to be sparse in time and, as the next theorem shows, it reduces $\mathcal{D}_{comp}$ to solving a linear program. The theorem's proof is in Appendix C.

Theorem 10. *For any $A, B \in \mathcal{S}$, $\mathcal{D}_{comp}(A, B)$ with $\|.\|$ equal to the matrix 1-norm can be computed (in polynomial time) by solving a linear program.*

The metric $\mathcal{D}_{comp}$ allows to consider a scaling factor α in front of the term $\sum_t \|W(t+1) - W(t)\|$, which we call the switch term, swi. We call the other term the distance term $dist$. Like the thr value in the CLEAR MOT, α allows us to penalize heavily the error in distances in some applications ($\alpha = small$); and in other applications allows us to penalize heavily the number of associations switches ($\alpha = large$). One can however report how close a set $A \in \mathcal{S}$ is from $B \in \mathcal{S}$ without choosing a specific α, simply plot the tradeoff curve between swi and $dist$ as a function of α. The closer this curve is to the left/bottom edges of the positive quadrant of the cartesian plane, the closer A to B.

The next theorem shows that the tradeoff curve produced by the two terms of $\mathcal{D}_{\text{comp}}$, which we denote by $(swi_{\mathcal{D}_{\text{comp}}}(\alpha), dist_{\mathcal{D}_{\text{comp}}}(\alpha))$, is, in a sense, the best trade-off curve between $dist$ and swi that any metric can generate. To state the theorem we need to define the region $\mathcal{R}$ of pairs of distance-value and switch-value that we can improve upon with some stochastic matrices $\{W(t)\}$ produced by some association procedure of some similarity measure. More concretely, $\mathcal{R} = \{(dist, swi) \in \mathbb{R}^2 : swi \geq \sum_t \|W(t+1) - W(t)\|$ and $dist \geq \sum_t \mathbf{tr}(W(t)^\dagger D^{AB}(t))$ for all $\{W(t)\} \in \mathcal{P}^T\}$. The proof of following result is in Appendix C.

Theorem 11. *If $(dis, swi) \in$ int$(\mathcal{R})$ then there exists α, such that $dis + \alpha swi > dis_{\mathcal{D}_{comp}}(\alpha) + \alpha swi_{\mathcal{D}_{comp}}(\alpha)$. If $(dis, swi) \in \partial\mathcal{R}$ then (dis, swi) is equal to some point in the curve $(dis_{\mathcal{D}_{comp}}(\alpha), swi_{\mathcal{D}_{comp}}(\alpha))$ for some α or is equal to a linear combination of points in the curve.*

Remark 3. *In the proof of the above theorem we show that $\mathcal{R}$ is convex. Hence, because the boundary of $\mathcal{R}$ is by definition Pareto efficient, this boundary must be a convex non-increasing curve. Since the curve $(dis_{\mathcal{D}_{comp}}(\alpha), swi_{\mathcal{D}_{comp}}(\alpha))$ effectively covers this boundary, we can use this property to efficiently estimate the trade-off curve. Basically, we can compute $(dis_{\mathcal{D}_{comp}}(\alpha), swi_{\mathcal{D}_{comp}}(\alpha))$ for some discrete set values of α and then we can take the convex hull of these points.*

A. Computing $\mathcal{D}_{comp}$ in practice

There are multiple ways to compute $\mathcal{D}_{comp}$ in practice, all of which amount to solving a convex optimization problem. For this paper we coded a solver in C based on the Alternating Direction Method of Multipliers (ADMM) because the ADMM is known to scale well for large problems and its modular nature makes it easy, in future work, for us to research variants of the formulation $\mathcal{D}_{comp}$ without having to re-write much of our code. For a good introduction to the ADMM, related methods and applications see [36]. Our implementation is approximately $\times 50$ faster than using Matlab CVX [37].

In Figure 2 we show the run-time of our code as a function of the total duration T of the input data for a different number of free association variables $W_{ij}(t)$ per time instant t. We use the terminology "free association variables" to emphasize that some $W_{ij}(t)$ might be a priory set to zero if the distance

between point i and point j at time t is very large, which can save computation time. In the figure we choose the euclidean distance for $d(.,.)$ and the one-norm for matrix norm $\|.\|$ in the switch term. Similar run-times hold for other metrics. The run-time is computed for randomly generated sets of trajectories on a single core of a 1.4GHz Intel Core i5 MacBook Air. In our tests, the ADMM always converged to 1% accuracy in less than 150 iterations.

To interpret the plots, imagine you want to evaluate the quality of a tracker when tracking 22 players in a soccer field. Imagine that the tracker operates at 30 frames per second and also that your tracker is noisy so that it produces a few false tracks that create approximately 10 extra points per instant. To compute $\mathcal{D}_{comp}$, and after you extend the ground-truth and hypothesis sets from A and B to A^+ and B^+, you are dealing with distance matrices with about $((22+10)\times 2)^2 = 4096$ variables per instant t. Using Fig. 2 you can conclude that it will take you about 40 seconds to evaluate the accuracy of $800/30 = 26.6$ seconds of data if you use a C implementation of the ADMM on a 1.4GHz MacBook Air. If you reduce the number of free variables per frame to half by setting $W_{ij}(t)$ to zero if point i and j are larger than a given threshold then you can reduce the time to process 26.6 seconds of data to about 20 seconds.

VI. More numerical results

To the best of our knowledge, $\mathcal{D}_{comp}$ is the first similarity measure that is a mathematical metric and deals with identity switches in an optimal way while being computable in polynomial time. In particular, we can use it to generate optimal tradeoff curves between $dist$ and swi (c.f. Theorem 11).

To illustrate this point we now use $\mathcal{D}_{comp}$ to assess the closeness between sets A and B both synthetically generated and coming from a tracking application and compare it with MOTP. We make the comparison through $dist$-swi tradeoff plots. For $\mathcal{D}_{comp}$ we use the extended Euclidean metric for d^+, the component-wise 1-norm for $\|.\|$ and the tradeoff curve is parametrized by a multiplying factor $\alpha > 0$ that we put in front of the switch term. For the MOTP, we compute the value for and swi in the same way as for $\mathcal{D}_{comp}$ but using the MOT heuristic association between A and B described in Definition 3. The MOTP tradeoff curve is parameterized by thr.

A direct interpretation of our results is that $\mathcal{D}_{comp}$ is better than MOTP. However, a more important interpretation of our results is that $\mathcal{D}_{comp}$ builds better associations between A and B than the greedy associations that are widely used in the literature and the non-greedy association that do not allow switches, like in OSPA. Because of this, most measures of accuracy, that are dependent on first establishing a correspondence between A and B, might be improved if we use the association that $\mathcal{D}_{comp}$ builds. Therefore, although we restrict our numerical comparison to MOTP, indeed it is one of the most widely used metrics in computer vision, we can use our work to also improve, for example, Multiple Object Tracking Accuracy (MOTA), False Alarms per Frame, Ratio of Mostly Tracked trajectories, Ratio of Mostly Lost trajectories, number of False Positives, number of False Negatives, number of ID Switches, number of tracks Fragmentation and many of the measures in [26], [27].

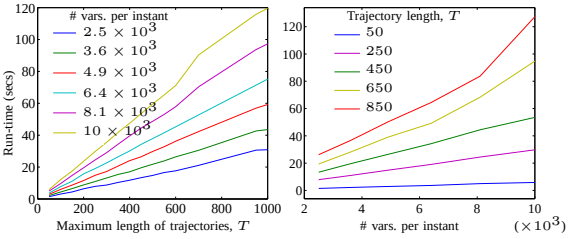

Fig. 2. Time it takes to compute $\mathcal{D}_{comp}$ (within 1% accuracy) as function of (a) the length of trajectories for a different number of association variables per instant; as a function of (b) the number of association variables per instant for different lengths of trajectories.

A. Real trackers and real data

In Figure 3-(a) we show the performance of the trackers in [38] and [39] on the AVG-TownCentre data set. We call these trackers $Tracker09$ and $Tracker11$ respectively. The data set is part of the Multiple Object Tracking Benchmark [40] that is widely used in computer vision. It comes from a pedestrian street filmed from an elevated point at a resolution of 1920×1080 for 3 minutes and 45 seconds and can be downloaded from [41]. In Figure 3-(b) we show the performance of the trackers in [42] and [43] on the PETS2009 data set, also part of the Multiple Object Tracking Benchmark. We call these trackers $Tracker12$ and $Tracker15$ respectively. Its resolution is 768×576, its duration is 1 minute and 54 seconds and it can be downloaded from [44]. We produced all the plots using the exact same output that each tracker produced in its respective paper thanks to the authors who provided us with their trackers' output. For computing we set $M = 20$ for AVG-Towncenter and $M = 50$ for PETS2009.

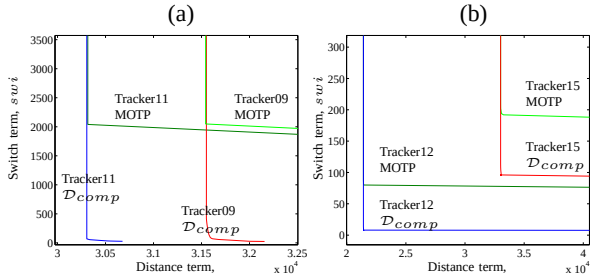

Fig. 3. (a) -swi tradeoff plot for $Tracker09$ and $Tracker11$ on the AVG-TownCenter data set; (b)dis-swi tradeoff plot for $Tracker12$ and $Tracker15$ on the PETS2009 data set. All the curves go to $swi = 0$ for large $dist$.

As expected from Theorem 11, the tradeoff curves from MOTP understate the performance of the trackers. In particular, all trackers are achieving substantial smaller number of switches without incurring in larger distance costs that what MOTP reports. Interestingly, for these trackers and datasets $\mathcal{D}_{comp}$ keeps the same relative ordering of performance as MOTP, although it is conceivable that there are situations in which a tracker 1 is better than a tracker 2 according to MOTP but not according to $\mathcal{D}_{comp}$. It would be interesting to find such an example. Finally notice that although the difference between the curves is large in absolute value, in relative values the curves are not very different. The absolute values are large because both and swi terms are sums of distances and switches

over all tracks and over all time instants and the relative values are small because all trackers are state-of-the-art and perform almost as well as possible.

B. Random ensemble of trajectories

In the examples above A and B are relatively close to each other because the trackers are all good trackers. In this section, to better appreciate the behavior of $\mathcal{D}_{comp}$, we randomly generate a set A with 25 trajectories (slightly more than e.g. the number of players in a football game) and then a set B of trajectories that are a distorted version of the trajectories in A. We generate the trajectories in A by randomly starting and ending a trajectory in time and having the object in that trajectory randomly change its velocity's direction along the way. The trajectories in B are generated by randomly fragmenting the trajectories in A, randomly removing some of the resulting trajectories, randomly adding noise to all trajectories and randomly flipping or not the ID of two trajectories if they pass by each other close enough. In the end, B might have more or less than 25 trajectories. In total we have four knobs to increase/reduce the distance between A and B: the amplitude of noise, $AMPnoise$, the probability of fragmenting a track at each point in time, $FRAGprob$, the probability of deleting a points in the track, $DELprob$, and the threshold distance after which we allow to tracks ID to be switched or not randomly, $SWIdist$.

These random trajectories are far more diverse than the trajectories in most publicly available data sets because real objects, like people, have fairly simple trajectories In addition, we do not just test two datasets, like in Section VI-A, but we datasets form about 20 different levels of distortion for each of the 4 knobs described above and for each of these levels of distortion we test 30 random sets A and B. Hence, if you will, and in the context of computer vision tracking, it is as if we test 2400 different data sets of ground-truth and hypothesis trajectories.

We study the similarity between the random sets A and B using $dist$-swi trade-off plots. The smaller the area under the curve of a the trade-off plot, the more similar A and B are. In Figure 4 we show the average area under the trade-off curve (AUC) for the random sets A and B under the different knob setting. The AUC is normalized by the largest AUC possible. The largest AUC is the product of the largest distance-error possible with the largest switch-error possible. Each point in the plots is an average over 30 random pairs A and B with the same knobs setting. In each plot we keep all but one knob constant.

As expected from Theorem 11, the AUC of $\mathcal{D}_{comp}$ is smaller than the AUC of MOTP. Note that it is incorrect to interpret these results as saying that MOTP is a like a scaled version of $\mathcal{D}_{comp}$. Recall that we use the same d^+ to measure the distance between points in space in both MOTP and $\mathcal{D}_{comp}$. We also compute the number of switches that the association of MOTP produces in the same way that we compute the number of switches that the association in $\mathcal{D}_{comp}$ produces. The difference in the curves of Figure 4 as a much deeper significance than a simple rescaling: for the 600 random pairs that we test, which consider different kinds of distortions

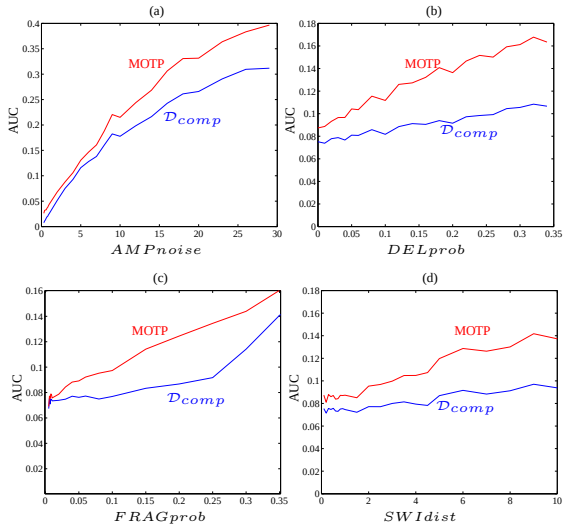

Fig. 4. AUC (normalized) versus (a) noise amplitude; (b) point deletion probability; (c) fragmenting probability; (d) switching distance. Smaller is better.

between A and B, and if we interpret each A and B as the ground-truth and output of a tracker, MOTP always says that the tracker is worse than it actually is. The $\mathcal{D}_{comp}$ can see similarities between A and B that the MOTP cannot.

VII. Conclusion

The problem of defining a similarity measure between sets of trajectories is crucial for computer vision, machine learning and general AI. An essential aspect in defining a similarity measure is finding a good and simple association between the elements of the sets. Existing similarity measures that define useful associations fail to be a metric mathematically speaking and the ones that are a metric only consider very primitive associations: full trajectories to full trajectories. The metric $\mathcal{D}_{comp}$ is the first that simultaneously (1) is a metric, (2) compares parts of trajectories to parts of trajectories, (3) allows to control the complexity of the association between trajectories in an globally optimal way and (4) can be computed in polynomial time.

The general idea of defining mathematical metrics for sets of trajectories using convex programs is our greatest overarching contribution. The next step is to explore variants of our metric that allow to incorporate uncertainty, as well as richer comparisons between A and B without losing its good properties.

References

[1] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *Signal Processing, IEEE Trans. on*, vol. 56, no. 8, pp. 3447–3457, 2008.

[2] B. Keni and S. Rainer, "Evaluating multiple object tracking performance: the clear mot metrics," *EURASIP Journal on Image and Video Processing*, vol. 2008, 2008.

[3] V. Ganti, R. Ramakrishnan, J. Gehrke, A. Powell, and J. French, "Clustering large datasets in arbitrary metric spaces," in *Data Engineering, 1999. Proceedings., 15th Intern. Conf. on*. IEEE, 1999, pp. 502–511.

[4] J. Kleinberg and E. Tardos, "Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields," *Journal of the ACM (JACM)*, vol. 49, no. 5, pp. 616–639, 2002.

[5] P. N. Yianilos, "Data structures and algorithms for nearest neighbor search in general metric spaces," in *Proceedings of the fourth annual ACM-SIAM Symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1993, pp. 311–321.

[6] B. Ristic, B.-N. Vo, D. Clark, and B.-T. Vo, "A metric for performance evaluation of multi-target tracking algorithms," *Signal Processing, IEEE Trans. on*, vol. 59, no. 7, pp. 3452–3457, 2011.

[7] B. E. Fridling and O. E. Drummond, "Performance evaluation methods for multiple-target-tracking algorithms," in *Orlando'91, Orlando, FL*. Intern. Society for Optics and Photonics, 1991, pp. 371–383.

[8] O. E. Drummond and B. E. Fridling, "Ambiguities in evaluating performance of multiple target tracking algorithms," in *Aerospace Sensing*. Intern. Society for Optics and Photonics, 1992, pp. 326–337.

[9] M. Betke, D. E. Hirsh, A. Bagchi, N. I. Hristov, N. C. Makris, and T. H. Kunz, "Tracking large variable numbers of objects in clutter," in *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conf. on*. IEEE, 2007, pp. 1–8.

[10] M. Betke, D. E. Hirsh, N. C. Makris, G. F. McCracken, M. Procopio, N. I. Hristov, S. Tang, A. Bagchi, J. D. Reichard, J. W. Horn *et al.*, "Thermal imaging reveals significantly smaller brazilian free-tailed bat colonies than previously estimated," *Journal of Mammalogy*, vol. 89, no. 1, pp. 18–24, 2008.

[11] Z. Wu, N. I. Hristov, T. L. Hedrick, T. H. Kunz, and M. Betke, "Tracking a large number of objects from multiple views," in *Computer Vision, 2009 IEEE 12th Intern. Conf. on*. IEEE, 2009, pp. 1546–1553.

[12] M. M. Deza and E. Deza, *Encyclopedia of distances*. Springer, 2009.

[13] O. Fujita, "Metrics based on average distance between sets," *Japan Journal of Industrial and Applied Mathematics*, vol. 30, no. 1, pp. 1–19, 2013.

[14] A. Gardner, J. Kanno, C. A. Duncan, and R. Selmic, "Measuring distance between unordered sets of different sizes," in *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conf. on*. IEEE, 2014, pp. 137–143.

[15] J. R. Hoffman and R. P. Mahler, "Multitarget miss distance via optimal assignment," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Trans. on*, vol. 34, no. 3, pp. 327–336, 2004.

[16] T. Vu and R. Evans, "A new performance metric for multiple target tracking based on optimal subpattern assignment," in *Information Fusion (FUSION), 2014 17th Intern. Conf. on*. IEEE, 2014, pp. 1–8.

[17] H. Xiaofan, R. Tharmarasa, T. Kirubarajan, and T. Thayaparan, "A track quality based metric for evaluating performance of multitarget filters," *Aerospace and Electronic Systems, IEEE Trans. on*, vol. 49, no. 1, pp. 610–616, 2013.

[18] S. Nagappa, D. E. Clark, and R. Mahler, "Incorporating track uncertainty into the ospa metric," in *Information Fusion (FUSION), 2011 Proceedings of the 14th Intern. Conf. on*. IEEE, 2011, pp. 1–8.

[19] T. Ellis, "Performance metrics and methods for tracking in surveillance," in *Proceedings of the 3rd IEEE Intern. Workshop on Performance Evaluation of Tracking and Surveillance (PETS?02)*. Citeseer, 2002, pp. 26–31.

[20] A. Milan, K. Schindler, and S. Roth, "Challenges of ground truth evaluation of multi-target tracking," in *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conf. on*. IEEE, 2013, pp. 735–742.

[21] O. E. Drummond, "Multiple-object estimation," *UCLA Ph.D. Dissertation*, 1975.

[22] S. B. Colegrove, L. Davis, and S. J. Davey, "Performance assessment of tracking systems," in *Signal Processing and Its Applications, 1996. ISSPA 96., Fourth Intern. Symposium on*, vol. 1. IEEE, 1996, pp. 188–191.

[23] F. Yin, D. Makris, and S. A. Velastin, "Performance evaluation of object tracking algorithms," in *10th IEEE Intern. Workshop on Performance Evaluation of Tracking and Surveillance (PETS2007), Rio de Janeiro, Brazil*, 2007.

[24] F. Bashir and F. Porikli, "Performance evaluation of object detection and tracking systems," in *IEEE Intern. Workshop on Performance Evaluation of Tracking and Surveillance (PETS)*, vol. 5, 2006.

[25] O. E. Drummond, "Methodologies for performance evaluation of multitarget multisensor tracking," in *SPIE's Intern. Symposium on Optical Science, Engineering, and Instrumentation*. Intern. Society for Optics and Photonics, 1999, pp. 355–369.

[26] R. L. Rothrock and O. E. Drummond, "Performance metrics for multiple-sensor multiple-target tracking," in *AeroSense 2000*. Intern. Society for Optics and Photonics, 2000, pp. 521–531.

[27] A. A. Gorji, R. Tharmarasa, and T. Kirubarajan, "Performance measures for multiple target tracking problems," in *Information Fusion (FUSION), 2011 Proceedings of the 14th Intern. Conf. on*. IEEE, 2011, pp. 1–8.

[28] S. Pingali and J. Segen, "Performance evaluation of people tracking systems," in *Applications of Computer Vision, 1996. WACV'96., Proceedings 3rd IEEE Workshop on*. IEEE, 1996, pp. 33–38.

[29] K. K. Edward, P. D. Matthew, and B. H. Michael, "An information theoretic approach for tracker performance evaluation," in *Computer Vision, 2009 IEEE 12th Intern. Conf. on*. IEEE, 2009, pp. 1523–1529.

[30] F. Porikli, "Trajectory distance metric using hidden markov model based representation," in *IEEE European Conf. on Computer Vision, PETS Workshop*, vol. 3, 2004.

[31] A. Cayley, "Lxxvii. note on the theory of permutations," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 34, no. 232, pp. 527–529, 1849.

[32] M. G. Kendall, "A new measure of rank correlation," *Biometrika*, pp. 81–93, 1938.

[33] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[34] M. R. Garey and D. S. Johnson, *Computers and intractability*. wh freeman, 2002, vol. 29.

[35] L. Khachiian, "Polynomial algorithm in linear programming," in *Akademiia Nauk SSSR, Doklady*, vol. 244, 1979, pp. 1093–1096.

[36] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[37] I. CVX Research, "CVX: Matlab software for disciplined convex programming, version 2.0," http://cvxr.com/cvx, Aug. 2012.

[38] B. Benfold and I. Reid, "Guiding visual surveillance by tracking human attention." in *BMVC*, 2009, pp. 1–11.

[39] ——, "Stable multi-target tracking in real-time surveillance video," in *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conf. on*. IEEE, 2011, pp. 3457–3464.

[40] "Multiple object tracking benchmark," http://www.motchallenge.net, accessed: 2015-03-01.

[41] "Coarse gaze estimation in visual surveillance," http://www.robots.ox.ac.uk/ActiveVision/Research/Projects/2009bbenfold_headpose/pro accessed: 2015-03-01.

[42] B. Yang and R. Nevatia, "Multi-target tracking by online learning of non-linear motion patterns and robust appearance models," in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conf. on*. IEEE, 2012, pp. 1918–1925.

[43] F. Poiesi and A. Cavallaro, "Tracking multiple high-density homogeneous targets," *Circuits and Systems for Video Technology, IEEE Trans. on*, vol. 25, no. 4, pp. 623–637, 2015.

[44] "Pets2009 benchmark data set," http://www.cvg.reading.ac.uk/PETS2009/a.html, accessed: 2015-03-01.

APPENDIX A
THE LIMITATIONS OF THE CLEAR MOT

Proof of Theorem 1. We will construct an example for which the theorem is true. Our example proves the theorem for any $1 < thr < 2$ and for $A = \{A_1, A_2\}$ and $B = \{B_1, B_2\}$, i.e. $m = 2$, where A_i and B_i are trajectories living in 1D. We explain how to generalize this example to any thr and any m at the end.

Consider two sets of one-dimensional trajectories $A = \{A_1, A_2\}$ and $B = \{B_1, B_2\}$ defined as in Figure 5 but where the trajectories extend beyond instant 15 until time T.

If $1 < thr < 2$ the CLEAR MOT builds a sequence Σ_{MOT} that is equal to $\{(1,2), ..., (1,2), (2,1), (2,1), ...\}$ because at some instant between time 4 and time 12 the initial association $\{A_1 \leftrightarrow B_1, A_2 \leftrightarrow B_2\}$ exceeds $thr < 2$ and is replaced by $\{A_2 \leftrightarrow B_1, A_1 \leftrightarrow B_2\}$ and after instant 12 this last association is anchored given that $1 < thr$.

The number of times that $\sigma_{\text{MOT}}(t) \neq \sigma_{\text{MOT}}(t+1)$ is 1 so $swi(\Sigma_{\text{MOT}}) = \frac{1}{T-1} = \mathcal{O}(1/T)$. After instant 12 we have $\sum_i \|A_i(t) - B_{\sigma_{\text{MOT}_i}(t)}(t)\| = 2$ so $dist(\Sigma_{\text{MOT}}, A, B) > \frac{2(T-12)}{T} = 2 - \mathcal{O}(1/T)$. However, if we choose $\Sigma =$

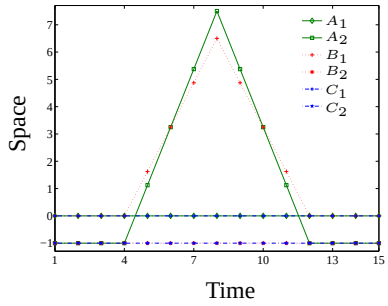

Fig. 5. Example that shows that (a) the CLEAR MOT association heuristic is bad and (b) MOTP does not define a metric.

$\{(1,2),...,(1,2)\}$ we have $swi(\Sigma) = 0$ and $dist(\Sigma, A, B) < \frac{12\times 7.5}{T} = \mathcal{O}(1/T)$.

We can make the proof hold for an interval around any thr by rescaling space. This changes the bounds on $dist$ by a factor of thr and leaves the bounds on swi unchanged. In addition, we can make the proof hold for any m (even) by extending A and B as follows. If the 1D trajectory A_i above is equal to $\{(1, A_i(1)),...,(T, A_i(T))\}$ define the 2D trajectory $A_i^{(k)}$ to have time-state pairs $\{(1,[A_i(1)); Ck],...,(T,[A_i(T); Ck])\}$ where C is a constant large enough to guarantee that trajectories for different ks are not close to each other. Define $B_i^{(k)}$ in a similar way. To make the proof hold for any m (even) extend A to $A = \{A_1^{(0)}, A_2^{(0)}, A_1^{(1)}, A_2^{(1)},..., A_1^{(m-1)}, A_2^{(m-1)}\}$ and extend B in a similar way. To extend the proof for an odd m simply append to A and B two equal trajectories far away from all other trajectories such that they are matched and hence contribute nothing to both swi or $dist$. With these sets the bounds previously computed on $swi()$ change by a factor of m and the bounds on $dist(.)$ change by a factor of $m/2$. In the statement of our theorem statement we are using $\mathcal{O}(m/T) = \mathcal{O}(1/T)$ and $\mathcal{O}(m\,thr/T) = \mathcal{O}(1/T)$. □

Proof of Theorem 2. We construct $A, B, C \in \mathcal{S}$ for which the triangle inequality is violated, specifically, $\mathcal{D}(A,B) > \mathcal{D}(A,C) + \mathcal{D}(C,B)$. Our example proves that the theorem holds for $1 < thr < 2$ but by rescaling space the proof follows for any $thr > 0$.

Consider the sets $A = \{A_1, A_2\}$, $B = \{B_1, B_2\}$ and $C = \{C_1, C_2\}$ as in Figure 5 but where the trajectories extend beyond time 15 to some large T. To make calculations simpler, we work with $\mathcal{D}$ divided by T in Definition 2.

Let us compute $\mathcal{D}(A,B)$ first. The association Σ_{MOT} for this distance is $\{(1,2),..,(1,2),(2,1),...\}$ because (1) we start with the association $\{A_1 \leftrightarrow B_1, A_2 \leftrightarrow B_2\}$, (2) at some point between instant 4 and instant 12 we need to change the association to $\{A_1 \leftrightarrow B_2, A_2 \leftrightarrow B_1\}$ because the initial association exceeds $thr < 2$ and (3) after time 12 the association $\{A_1 \leftrightarrow B_2, A_2 \leftrightarrow B_1\}$ is anchored because $thr > 1$. This association leads to $\mathcal{D}(A,B) > \frac{2(T-12)}{T}$.

Now we compute $\mathcal{D}(A,C)$. The computation for $\mathcal{D}(C,B)$ is similar so we omit it. The association for $\mathcal{D}(A,C)$ is $\Sigma_{\text{MOT}} = \{(1,2),...,(1,2)\}$ because (1) we start with $\{A_1 \leftrightarrow C_1, A_2 \leftrightarrow C_2\}$, (2) the association $A_1 \leftrightarrow C_1$ is always anchored because the distance between A_1 and C_1 is always zero and thus always smaller than $thr > 1$ and (3) after some point between instant 4 and instant 12, when the distance between A_2 and C_2 exceeds $thr < 2$, MOTP still keeps the association $A_2 \leftrightarrow C_2$ because A_1 and C_1 are already anchored. This association leads to $\mathcal{D}(A,C) < \frac{8.5}{T}$.

Therefore, for T large enough we have $\mathcal{D}(A,B) > \frac{2(T-12)}{T} > \frac{8.5}{T} + \frac{8.5}{T} > \mathcal{D}(A,C) + \mathcal{D}(C,B)$. □

APPENDIX B
A NATURAL METRIC FOR SETS OF TRAJECTORIES

To prove Theorem 3 we need the following Lemma.

Lema 2. *The map d^+ is a metric on $\mathbb{R}^p \cup \{*\}$.*

Proof. Let $x'', x', x \in \mathbb{R}^p \cup \{*\}$. To prove the coincidence property observe that $d^+(x,x') = 0$ either implies that $x = x' = *$ or, since $M > 0$, implies that $d^+(x,x') = d(x,x') = 0$, which, because d is a metric, implies that $x = x' = 0$. In other words, $d^+(x,x') = 0 \Leftrightarrow x = x'$. The symmetry property, $d^+(x,x') = d^+(x',x)$, follows directly from the definition. To prove the subadditivity property we need to consider eight different cases. It is trivial to check that the following six cases satisfy the triangle inequality: $x, x', x'' \in \mathbb{R}^p$; $x = x' = *, x'' \in \mathbb{R}^p$; $x'' = x' = *, x \in \mathbb{R}^p$; $x = *, x', x'' \in \mathbb{R}^p$; $x'' = *, x, x' \in \mathbb{R}^p$; $x = x'' = *, x' \in \mathbb{R}^p$. We check the two other cases separately. If $x, x', x'' \in \mathbb{R}^p$ then $d^+(x,x'') = \min\{2M, d(x,x'')\} \leq \min\{2M, d(x,x') + d(x',x'')\} \leq \min\{2M, d(x,x')\} + \min\{2M, d(x',x'')\} = d^+(x,x') + d^+(x',x'')$. If $x' = *, x, x'' \in \mathbb{R}^p$ then $d^(x,x'') = \min\{2M, d(x,x'')\} \leq M + M = d^+(x,x') + d^+(x',x'')$. □

Proof of Theorem 3. Let A, B, C be any three elements in $\mathcal{S}$. *Coincidence property*: We show that $\mathcal{D}_{nat}(A,B) = 0$ if and only if $A = B$. Remember that A and B are unordered sets of trajectories so $A = B$ means that there is an isomorphism between A and B. In other words, they are equal apart tom a relabeling of their elements. If $A = B$ and we set $\Sigma = (\sigma, \sigma, ..., \sigma)$ where σ is an isomorphism between A and B, then the objective in (2) is equal to zero. Since the minimum of (2) must always be non-negative, we conclude that $A = B \Rightarrow \mathcal{D}_{nat}(A,B) = 0$. Now assume that $\mathcal{D}_{nat}(A,B) = 0$ and let $\Sigma^* = (\sigma^*(1),...,\sigma^*(T))$ be a minimizer of (2). $\mathcal{D}_{nat}(A,B) = 0$ implies that $\mathcal{K}(\Sigma^*) = 0$ and therefore $\sigma_i^*(t) = \sigma_i^*(1)$ for all t and i. Since the labeling of the trajectories does not matter in computing $\mathcal{D}_{nat}$, we assume without loss of generality that their labeling is such that we can write $\sigma_i^*(t) = i$. $\mathcal{D}_{nat}(A,B) = 0$ also implies that, for all t and i, we have $d^+(A_i^+(t), B^+_{\sigma_i^*(t)}(t)) = d^+(A_i^+(t), B_i^+(t)) = 0$. Since d^+ is a metric this in turn implies that $A_i^+(t) = B_i^+(t)$ for all i and t, which is the same as saying that $A^+ = B^+$. If $A \neq B$ then $A^+ \neq B^+$ therefore $A^+ = B^+$ implies that $A = B$. To be more specific, A is equal to B apart from a relabeling of its trajectories, which we can do because A and B are unordered sets of trajectories.
Symmetry property: $\mathcal{D}_{nat}$ only depends on A and B through d^+ and, like all metrics, this is a symmetric function. In addition, if we swap i and $\sigma_i(t)$ in (2) the minimum of (2) remains unchanged. It follows that $\mathcal{D}_{nat}(A,B) = \mathcal{D}_{nat}(B,A)$.
Subadditivity property: We prove that $\mathcal{D}_{nat}(A,C) \leq$

$\mathcal{D}_{nat}(A,B)+\mathcal{D}_{nat}(B,C)$. Since d^+ is a metric, we can write that

$$d^+(A_i^+(t),C^+_{\sigma_i(t)}(t)) \leq d^+(A_i^+(t),B^+_{\sigma'_{\sigma_i(t)}(t)}(t)) + d^+(B^+_{\sigma'_{\sigma_i(t)}(t)}(t),C^+_{\sigma_i(t)}(t)) \quad (5)$$

for any $\Sigma'=(\sigma'(1),...,\sigma'(T))\in\Pi^T$ and for all i and t. Now notice that

$$\sum_{i=1}^m d^+(B^+_{\sigma'_{\sigma_i(t)}(t)}(t),C^+_{\sigma_i(t)}(t)) = \sum_{i=1}^m d^+(B^+_{\sigma'_i(t)}(t),C^+_i(t)).$$

Using this together with (5) we can write

$$\sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),C^+_{\sigma_i(t)}(t)) \leq \sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),B^+_{\sigma'_{\sigma_i(t)}(t)}(t)) + \sum_{t=1}^T\sum_{i=1}^m d^+(B^+_{\sigma'_i(t)}(t),C^+_i(t)).$$

Let us define $\Sigma''=(\sigma''(1),...,\sigma''(T))\in\Pi^T$ where $\sigma''_i(t)=\sigma'_{\sigma_i(t)}(t)$. This means that $\Sigma''=\Sigma'\circ\Sigma$. We can use Σ' to rewrite the above expression as

$$\sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),C^+_{\sigma_i(t)}(t)) \leq \sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),B^+_{\sigma''_i(t)}(t)) + \sum_{t=1}^T\sum_{i=1}^m d^+(B^+_{\sigma'_i(t)}(t),C^+_i(t)). \quad (6)$$

If we use the second and third property that $\mathcal{K}$ satisfies we can also write

$$\mathcal{K}(\Sigma)=\mathcal{K}(\Sigma'^{-1}\circ\Sigma'\circ\Sigma)\leq\mathcal{K}(\Sigma'^{-1})+\mathcal{K}(\Sigma'\circ\Sigma) = \mathcal{K}(\Sigma')+\mathcal{K}(\Sigma''). \quad (7)$$

Now we add both sides of (6) and (7) and obtain $\mathcal{K}(\Sigma)+\sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),C^+_{\sigma_i(t)}(t)) \leq \mathcal{K}(\Sigma'')+\sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),B^+_{\sigma''_i(t)}(t))+\mathcal{K}(\Sigma')+\sum_{t=1}^T\sum_{i=1}^m d^+(B^+_{\sigma'_i(t)}(t),C^+_i(t))$. Finally, we find the minimum of both sides of the inequality over all pairs of Σ and Σ'. Recall that we can choose Σ' independently of Σ. Since for every pair (Σ,Σ') there is exactly one pair (Σ',Σ''), finding the minimum over Σ and Σ' is the same as finding the minimum over Σ' and Σ''. Therefore we get,

$$\min_{\Sigma\in\Pi^T}\mathcal{K}(\Sigma)+\sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),C^+_{\sigma_i(t)}(t)) \leq \min_{\Sigma''\in\Pi^T}\mathcal{K}(\Sigma'')+\sum_{t=1}^T\sum_{i=1}^m d^+(A_i^+(t),B^+_{\sigma''_i(t)}(t)) + \min_{\Sigma'\in\Pi^T}\mathcal{K}(\Sigma')+\sum_{t=1}^T\sum_{i=1}^m d^+(B^+_{\sigma'_i(t)}(t),C^+_i(t)),$$

which is the same as $\mathcal{D}_{nat}(A,C)\leq\mathcal{D}_{nat}(A,B)+\mathcal{D}_{nat}(B,C)$. □

Proof of Theorem 4. The first two properties are immediate to check. To prove the third property it sufficient to prove that $\mathbb{I}((\sigma'(t+1)\circ\sigma(t+1))\circ(\sigma'(t)\circ\sigma(t))^{-1}\neq I)\leq\mathbb{I}(\sigma(t+1)\circ\sigma(t)^{-1}\neq I)+\mathbb{I}(\sigma'(t+1)\circ\sigma'(t)^{-1}\neq I)$. To prove this we consider different cases. If $\sigma'(t+1)=\sigma'(t)$ and $\sigma(t+1)=\sigma(t)$ then the inequality holds as $0\leq 0+0$. If $\sigma'(t+1)\neq\sigma'(t)$ and $\sigma(t+1)\neq\sigma(t)$ then the inequality holds since the left hand side is at most 1 and the right hand side in this case is equal to $1+1$. If $\sigma(t+1)=\sigma(t)$ and $\sigma'(t+1)\neq\sigma'(t)$ then the inequality holds as $1\leq 0+1$. The final case is when $\sigma(t+1)\neq\sigma(t)$ and $\sigma'(t+1)=\sigma'(t)$. Notice that $\mathbb{I}((\sigma'(t+1)\circ\sigma(t+1))\circ(\sigma'(t)\circ\sigma(t))^{-1}\neq I)=\mathbb{I}(\sigma(t+1)\circ\sigma(t)^{-1}\neq\sigma'(t+1)^{-1}\circ\sigma'(t))$ and conclude that the inequality holds with $1\leq 1+0$. □

Proof of Theorem 5 and Theorem 6. We prove the theorem for $\mathcal{K}_{trans}$. The proof for $\mathcal{K}_{adjtrans}$ is very similar. The first two properties are immediate to check from the definition of $\mathcal{K}_{trans}$. To prove the third property it suffices to show that $k_{Cayley}((\sigma'(t+1)\circ\sigma(t+1))\circ(\sigma'(t)\circ\sigma(t))^{-1})\leq k_{Cayley}(\sigma(t+1))\circ\sigma(t)^{-1})+k_{Cayley}(\sigma'(t+1)\circ\sigma'(t))^{-1})$. To prove this we use of the following two facts. Fact 1: if $\sigma,\sigma'\in\Pi$ then $k_{Cayley}(\sigma\circ\sigma')\leq k_{Cayley}(\sigma)+k_{Cayley}(\sigma')$. Fact 2: if $\sigma,\sigma'\in\Pi$ then $k_{Cayley}(\sigma\circ\sigma'\circ\sigma^{-1})=k_{Cayley}(\sigma')$. Both facts are obvious if stated in plain english. Fact 1: The shortest set of transpositions that takes $\sigma\circ\sigma'$ to I is shorter than any set of transpositions that first takes $\sigma\circ\sigma'$ to σ' and then takes σ' to I. Fact 2: The shortest set of transpositions that takes a permutation of objects to the identity permutation is as short as the set of transpositions that take a relabeling of those objects to the identity permutation.

Let $\sigma_A=\sigma(t)\circ\sigma(t+1)^{-1}$, $\sigma_B=\sigma'(t)\circ\sigma'(t+1)^{-1}$ and let $\sigma_C=\sigma'(t)\circ\sigma_A\circ\sigma'(t)^{-1}$. Observe that the permutation $\sigma_C\circ\sigma_B$ satisfies

$$\sigma_C\circ\sigma_B\circ((\sigma'(t+1)\circ\sigma(t+1))\circ(\sigma'(t)\circ\sigma(t))^{-1})=I.$$

Therefore we can apply Fact 2 followed by Fact 1 and obtain $k_{Cayley}((\sigma'(t+1)\circ\sigma(t+1))\circ(\sigma'(t)\circ\sigma(t))^{-1})=k_{Cayley}(\sigma_B^{-1}\circ\sigma_C^{-1})\leq k_{Cayley}(\sigma_B^{-1})+k_{Cayley}(\sigma_C^{-1})=k_{Cayley}(\sigma_B^{-1})+k_{Cayley}(\sigma_A^{-1})=k_{Cayley}(\sigma(t+1)\circ\sigma(t)^{-1})+k_{Cayley}(\sigma'(t+1)\circ\sigma'(t)^{-1})$. □

Proof of Theorem 7. We give a counter example that violates property (iii) for $\beta=1$. It is easy to come up with similar counter examples that violate property (iii) for any value of $\beta\geq 1$. Let $I=(1,2)$ be the identity permutation and let $\sigma_0=(2,1)$ be the permutation that swaps 1 and 2. Let $\Sigma=(I,\sigma_0,\sigma_0)$ and let $\Sigma'=(I,I,\sigma_0)$. We have $\mathcal{K}_{maxcount}(\Sigma)=1$ and $\mathcal{K}_{maxcount}(\Sigma')=1$ but $\mathcal{K}_{maxcount}(\Sigma'\circ\Sigma)=\infty>\mathcal{K}_{maxcount}(\Sigma')+\mathcal{K}_{maxcount}(\Sigma)$. □

Proof of Theorem 8. We produce three sets of trajectories A, B and C that make (2) violate the triangle inequality. In this proof we assume $\beta=1$ but it is easy to change A, B and C such that the proof holds for any β,[1]. Let $I=(1,2)$ be the identity permutation and let $\sigma_0=(2,1)$ be the permutation that swaps 1 and 2. Let $A=\{A_1,A_2\}$ where $A_1=\{(1,2),(2,-2),(3,-2)\}$ and $A_2=\{(1,-2),(2,2),(3,2)\}$, let $B=\{B_1,B_2\}$ where $B_1=\{(1,2),(2,2),(3,2)\}$ and

[1]It is also not hard to see that we do not really need d to be the Euclidean distance for the proof to hold.

$B_2 = \{(1,-2),(2,-2),(3,-2)\}$ and let $C = \{C_1, C_2\}$ where $C_1 = \{(1,2),(2,2),(3,-2)\}$ and $C_2 = \{(1,-2),\ (2,-2),\ (3,2)\}$. Now consider (2) with $\mathcal{K} = \mathcal{K}_{maxcount}$ and $\beta = 1$. The optimization problem $\mathcal{D}_{nat}(A,B)$ has a minimum of 1 at $\Sigma = (I, \sigma_0, \sigma_0)$. The optimization problem $\mathcal{D}_{nat}(B,C)$ has a minimum of 1 at $\Sigma = (I, I, \sigma_0)$. When we solve the optimization problem $\mathcal{D}_{nat}(A,C)$ we are only allowed to perform one change in the association between A and C, otherwise the term $\mathcal{K}_{maxcount}$ makes us pay a very large cost. With only one change in the association, we incur in a distance of 4 for $t = 1$ or $t = 2$ or $t = 3$. Hence, $\mathcal{D}_{nat}(A,C) \geq 4 > 1 + 1 = \mathcal{D}_{nat}(A,B) + \mathcal{D}_{nat}(B,C)$. □

APPENDIX C
A NATURAL AND COMPUTABLE METRIC

Proof of Theorem 9. Let A, B and C be any three elements of $\mathcal{S}$.

Coincidence property: If $\mathcal{D}_{comp}(A,B) = 0$ then $W(t) = W(1)$ for all t. Hence, if $\mathcal{D}_{comp}(A,B) = 0$ then $D^{AB}(t)_{ij} = 0$ for all i, j such that $W(1)_{ij} > 0$. Now recall that if W_{ij} is a doubly stochastic matrix then there exists a permutation $\sigma \in \Pi$ such that $W_{i\sigma_i} > 0$ for all i. Therefore, there exist $\sigma \in \Pi$ such that $D^{AB}(t)_{i\sigma_i} = 0$ for all t and i. In other words, A and B are the same apart from a relabeling of their elements.

Symmetry property: Using the properties of **tr** we have that $\mathbf{tr}(W^\dagger(t) D^{AB}(t)) = \mathbf{tr}((W^\dagger(t)\ D^{AB}(t))^\dagger) = \mathbf{tr}(D^{AB}(t)^\dagger W(t)) = \mathbf{tr}(D^{BA}(t) W(t)) = \mathbf{tr}(W(t) D^{BA}(t))$. Minimizing with respect to $\{W(t)\}$ is the same as minimizing with respect to $\{W(t)^\dagger\}$ so $\mathcal{D}_{comp}(A,B) = \mathcal{D}_{comp}(B,A)$.

Subadditivity property: We prove that $\mathcal{D}_{comp}(A,C) \leq \mathcal{D}_{comp}(A,B) + \mathcal{D}_{comp}(B,C)$. First note that, since d^+ is a metric we have that $D^{AC}(t)_{ij} \leq D^{AB}(t)_{ik} + D^{BC}(t)_{kj}$ for any k. Let $\{W_1(t)\}, \{W_2(t)\} \in \mathcal{P}^T$. We multiply both sides of the previous inequality by $W_1(t)_{ik} W_2(t)_{kj}$ and sum over ijk and obtain $\sum_{i,j,k=1}^m W_1(t)_{ik} W_2(t)_{kj} D^{AC}(t)_{ij} \leq \sum_{i,j,k=1}^m W_1(t)_{ik} W_2(t)_{kj} D^{AB}(t)_{ik}$ $+ \sum_{i,j,k=1}^m W_1(t)_{ik} W_2(t)_{kj} D^{BC}(t)_{kj}$. Note that $W_1(t)$ and $W_2(t)$ are doubly stochastic matrices therefore we have $\sum_{j=1}^m W_2(t)_{kj} = 1$ and $\sum_{i=1}^m W_1(t)_{ik} = 1$ and we can re-write the previous inequality in matrix for as $\mathbf{tr}((W_1(t)W_2(t))^\dagger D^{AC}(t)) \leq \mathbf{tr}(W_1(t)^\dagger D^{AB}(t)) + \mathbf{tr}(W_2(t)^\dagger D^{BC}(t))$. From the definition of $\|\|$ we have that $\|W_1(t+1)W_2(t+1) - W_1(t)W_2(t)\| \leq \|W_1(t+1) - W_1(t)\| + \|W_2(t+1) - W_2(t)\|$. Summing both sides of the last two inequalities and, summing over t, and minimizing over $\{W_1(t)\}, \{W_2(t)\} \in \mathcal{P}^T$ we have that

$$\min_{\{W_1(t)\},\{W_2(t)\}\in\mathcal{P}^T} \sum_{t=1}^{T-1} \|W_1(t+1)W_2(t+1) - W_1(t)W_2(t)\| + \sum_{t=1}^{T} \mathbf{tr}((W_1(t)W_2(t))^\dagger D^{AC}(t))$$
$$\leq \min_{\{W_1(t)\}\in\mathcal{P}^T} \sum_{t=1}^{T-1} \|W_1(t+1) - W_1(t)\| + \sum_{t=1}^{T} \mathbf{tr}(W_1(t)^\dagger D^{AB}(t))$$
$$+ \min_{\{W_2(t)\}\in\mathcal{P}^T} \sum_{t=1}^{T-1} \|W_2(t+1) - W_2(t)\| + \sum_{t=1}^{T} \mathbf{tr}(W_2(t)^\dagger D^{BC}(t)).$$

Note that for any $W_1(t), W_2(t) \in \mathcal{P}$ we have that $W_1(t)W_2(t) \in \mathcal{P}$ so in the minimization performed on the left hand side we can replace $W_1(t)W_2(t)$ by a single $W(t)$. The subadditivity property follows. □

Proof of Lemma 1. Let $X_1, X_2, Y_1, Y_2 \in \mathcal{P}$. By an application of the triangle inequality followed by two applications of the sub-multiplicative property followed by two applications of the norm-bound property we have that $\|Y_2X_2 - Y_1X_1\| \leq \|Y_2(X_2 - X_1 + X_1) - Y_1X_1\| \leq \|Y_2(X_2 - X_1)\| + \|Y_2X_1 - Y_1X_1\| = \|Y_2(X_2 - X_1)\| + \|(Y_2 - Y_1)X_1\| \leq \|Y_2\|\|(X_2 - X_1)\| + \|(Y_2 - Y_1)\|\|X_1\| \leq \|(X_2 - X_1)\| + \|(Y_2 - Y_1)\|$. □

Proof of Theorem 10. Let us look at the optimization problem (4) more closely. First notice that the constraints that define $\mathcal{P}$ are a set of linear constraints. Second notice that the first term in the objective is a linear function of $\{W(t)\}$. Third notice that we can replace the term $\sum_{t=1}^{T-1} \|W(t+1) - W(t)\|$ in the objective by $\sum_{t=1}^{T-1} e_t$ if we add the additional constraints that $\|W(t+1) - W(t)\| \leq e_t$ for all t. Now we show that these additional constraints can be represented by linear constraints. Since $\|W(t+1) - W(t)\| = \max_j \sum_i |W(t+1)_{ij} - W(t)_{ij}|$ each of these constraints can be replaced by $\sum_i h(t)_{ij} \leq e_t$ for all t and j if we add the additional constrain that $|W(t+1)_{ij} - W(t)_{ij}| \leq h(t)_{ij}$ for all i, j and t. Each of these constrains can be written as $W(t+1)_{ij} - W(t)_{ij} \leq h(t)_{ij}$ and $-W(t+1)_{ij} + W(t)_{ij} \leq h(t)_{ij}$ and both of these constraints are linear so we are done. □

Proof of Theorem 11. We first prove that $\mathcal{R}$ is convex. Let $(a,b) \in \mathcal{R}$ and $(c,d) \in \mathcal{R}$. We have that $a \geq \sum_t \|W_1(t+1) - W_1(t)\|$ and $b \geq \sum_t \mathbf{tr}(W_1(t)^\dagger D^{AB}(t))$ and $c \geq \sum_t \|W_2(t+1) - W_2(t)\|$ and $d \geq \sum_t \mathbf{tr}(W_2(t)^\dagger D^{AB}(t))$ for some $\{W_1(t)\}, \{W_2(t)\} \in \mathcal{P}^T$. Let $W_3(t) = \beta W_1(t) + (1-\beta) W_2(t)$ for some $\beta \in [0,1]$ and all t. Since $\mathcal{P}^T$ is a convex set we have that $W_3(t) \in \mathcal{P}^T$. In addition, since the functions $\sum_t \|W(t+1) - W(t)\|$ and $\sum_t \mathbf{tr}(W(t)^\dagger D^{AB}(t))$ are convex functions of $\{W(t)\}$ we have that $\beta a + (1-\beta)c \geq \sum_t \|W_3(t+1) - W_3(t)\|$ and $\beta b + (1-\beta)d \geq \sum_t \mathbf{tr}(W_3(t)^\dagger D^{AB}(t))$. Therefore, we conclude that $\beta(a,b) + (1-\beta)(c,d) \in \mathcal{R}$.

We now prove the rest of the theorem. Since $\mathcal{R}$ is convex, we can write it as the following intersection of half-planes $\mathcal{R} = \{(x,y) \in \mathbb{R}^2 : x + \alpha y \geq \min_{\{W(t)\}\in\mathcal{P}^T} \alpha u + v \text{ s.t. } u \geq \sum_t \|W(t+1) - W(t)\| \text{ and } v \geq \sum_t \mathbf{tr}(W(t)^\dagger D^{AB}(t)) \text{ for all } \alpha \in [0,\infty]\}$. From this representation it directly follows that, for any $\alpha \in [0,\infty]$, any point $(dis, swi) \in \text{int}(\mathcal{R})$ satisfies $\alpha swi + dis >$ minimum $\{\alpha u + v$ subject to $W(t) \in \mathcal{P}^T; u \geq \sum_t \|W(t+1) - W(t)\|; v \geq \sum_t \mathbf{tr}(W(t)^\dagger D^{AB}(t))\} = \alpha \sum_t \|W^*(t+1) - W^*(t)\| + \sum_t \mathbf{tr}(W^*(t)^\dagger D^{AB}(t)) = \alpha swi_{\mathcal{D}_{\text{comp}}}(\alpha) + dis_{\mathcal{D}_{\text{comp}}}(\alpha)$, where $\{W^*(t)\} \in \mathcal{P}^T$ is the optimum point of the optimization problem of $\mathcal{D}_{comp}$.

From this representation we also conclude that any point at the boundary of $\mathcal{R}$ is either of the form $(\sum_t \|W^*(t+1) - W^*(t)\|, \sum_t \mathbf{tr}(W^*(t)^\dagger D^{AB}(t))) = (dis_{\mathcal{D}_{\text{comp}}}(\alpha), swi_{\mathcal{D}_{\text{comp}}}(\alpha))$ for some optimal $\{W^*(t)\} \in \mathcal{P}^T$ and some α or it is a convex combination of points of this form. □